

Freeform Search

Database:

- US Pre-Grant Publication Full-Text Database
- US Patents Full-Text Database
- US OCR Full-Text Database
- EPO Abstracts Database
- JPO Abstracts Database
- Derwent World Patents Index
- IBM Technical Disclosure Bulletins

Term:

Display: Documents in Display Format: Starting with Number **Generate:** Hit List Hit Count Side by Side Image

Search History

DATE: Saturday, March 05, 2005 [Printable Copy](#) [Create Case](#)

<u>Set Name</u>	<u>Query</u>	<u>Hit Count</u>	<u>Set Name</u>
side by side result set			

DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR

<u>L45</u>	L44 and servers	28	<u>L45</u>
<u>L44</u>	L43 and workstations	30	<u>L44</u>
<u>L43</u>	I34 and (I/O driver or on/off driver)	141	<u>L43</u>

DB=USPT; PLUR=YES; OP=OR

<u>L42</u>	5828876.pn.	1	<u>L42</u>
<u>L41</u>	5828876.pn.	1	<u>L41</u>
<u>L40</u>	5832222.pn.	1	<u>L40</u>
<u>L39</u>	5832222.pn.	1	<u>L39</u>
<u>L38</u>	5889520.pn.	1	<u>L38</u>
<u>L37</u>	5893086.pn.	1	<u>L37</u>

DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR

<u>L36</u>	L34 and (policy or polic\$) and protocol	31	<u>L36</u>
<u>L35</u>	L34 and (policy or polic\$) near protocol	3	<u>L35</u>
<u>L34</u>	L32 and stor\$ near node	153	<u>L34</u>
<u>L33</u>	L32 and shared near stor\$ near node	8	<u>L33</u>
<u>L32</u>	L31 and (network or internet or www)	2006	<u>L32</u>

Best Available Copy

h e b b g e e e e h f f f e ch

<u>L31</u>	(mirror\$ or shadow\$) near data	3677	<u>L31</u>
<u>L30</u>	370/417	483	<u>L30</u>
<u>L29</u>	370.clas.	82251	<u>L29</u>
<u>L28</u>	714/6	2652	<u>L28</u>
<u>L27</u>	714/4	2437	<u>L27</u>
<u>L26</u>	714.clas.	48219	<u>L26</u>
<u>L25</u>	707/104.1	5274	<u>L25</u>
<u>L24</u>	707/10	9752	<u>L24</u>
<u>L23</u>	707/9	2423	<u>L23</u>
<u>L22</u>	707/8	2188	<u>L22</u>
<u>L21</u>	707/1	6938	<u>L21</u>
<u>L20</u>	707.clas.	25375	<u>L20</u>
<u>L19</u>	711/147	1929	<u>L19</u>
<u>L18</u>	711/114	2648	<u>L18</u>
<u>L17</u>	711/111	1436	<u>L17</u>
<u>L16</u>	711.clas.	24509	<u>L16</u>
<u>L15</u>	709/232	2287	<u>L15</u>
<u>L14</u>	709/229	4553	<u>L14</u>
<u>L13</u>	709/227	4791	<u>L13</u>
<u>L12</u>	709/224	6818	<u>L12</u>
<u>L11</u>	709/212	730	<u>L11</u>
<u>L10</u>	709/217	6435	<u>L10</u>
<u>L9</u>	709/216	966	<u>L9</u>
<u>L8</u>	709/200	3343	<u>L8</u>
<u>L7</u>	709/201	4191	<u>L7</u>
<u>L6</u>	709/204	2480	<u>L6</u>
<u>L5</u>	709/203	8897	<u>L5</u>
<u>L4</u>	709/214	666	<u>L4</u>
<u>L3</u>	709/213	1947	<u>L3</u>
<u>L2</u>	709/219	5999	<u>L2</u>
<u>L1</u>	709.clas.	34731	<u>L1</u>

END OF SEARCH HISTORY

Freeform Search

Database:	US Pre-Grant Publication Full-Text Database US Patents Full-Text Database US OCR Full-Text Database EPO Abstracts Database JPO Abstracts Database Derwent World Patents Index IBM Technical Disclosure Bulletins
Term: <input style="width: 80%; height: 40px;" type="text"/> <div style="position: absolute; right: -10px; top: 0; width: 10px; height: 40px; background-color: black; opacity: 0.5;"></div>	
Display: <input type="text" value="10"/> Documents in <u>Display Format:</u> <input type="text"/> Starting with Number <input type="text" value="1"/>	
Generate: <input type="radio"/> Hit List <input checked="" type="radio"/> Hit Count <input type="radio"/> Side by Side <input type="radio"/> Image	

| |

Search History

DATE: Saturday, March 05, 2005 [Printable Copy](#) [Create Case](#)

<u>Set</u> <u>Name</u> <u>Query</u>	<u>Hit</u> <u>Count</u>	<u>Set</u> <u>Name</u>
side by side		result set
<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR</i>		
<u>L19</u> L18 and (mirror\$ or shadow\$)	68	<u>L19</u>
<u>L18</u> l1 and backup and restore	109	<u>L18</u>
<i>DB=USPT; PLUR=YES; OP=OR</i>		
<u>L17</u> 5519844.pn.	1	<u>L17</u>
<u>L16</u> 5519844.pn.	1	<u>L16</u>
<u>L15</u> 5659704.pn.	1	<u>L15</u>
<u>L14</u> 5659704.pn.	1	<u>L14</u>
<u>L13</u> 4414644.pn.	1	<u>L13</u>
<u>L12</u> 4533995.pn.	1	<u>L12</u>
<u>L11</u> 5673382.pn.	1	<u>L11</u>
<u>L10</u> 5673382.pn.	1	<u>L10</u>
<u>L9</u> 5829046.pn.	1	<u>L9</u>
<u>L8</u> 5857208.pn.	1	<u>L8</u>
<u>L7</u> 5857208.pn.	1	<u>L7</u>
<u>L6</u> 5875478.pn.	1	<u>L6</u>

<u>L5</u>	6023281.pn.	1	<u>L5</u>
<u>L4</u>	6076148.pn.	1	<u>L4</u>
<u>L3</u>	5991793.pn.	1	<u>L3</u>

DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=OR

<u>L2</u>	L1 and (mirror\$ or shadow\$) near data	91	<u>L2</u>
<u>L1</u>	virtual near (stor\$ or storage) and area and (network or www or internet)	1486	<u>L1</u>

END OF SEARCH HISTORY

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
 [Generate Collection](#) [Print](#)

L34: Entry 149 of 153

File: USPT

Oct 14, 1997

US-PAT-NO: 5678061

DOCUMENT-IDENTIFIER: US 5678061 A

TITLE: Method for employing doubly striped mirroring of data and reassigning data streams scheduled to be supplied by failed disk to respective ones of remaining disks

DATE-ISSUED: October 14, 1997

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Mourad; Antoine N.	Aberdeen	NJ		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Lucent Technologies Inc.	Murray Hill	NJ			02

APPL-NO: 08/ 504096 [PALM]

DATE FILED: July 19, 1995

INT-CL: [06] G06 F 15/02

US-CL-ISSUED: 395/841; 395/182.04, 395/405

US-CL-CURRENT: 710/21; 711/5, 714/6

FIELD-OF-SEARCH: 395/182.04, 395/182.05, 395/441, 395/439, 395/489, 395/841, 395/405, 395/497.04, 395/800, 395/842, 395/182.03, 371/40.4

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<input type="checkbox"/> <u>4688168</u>	August 1987	Gudaitis et al.	395/287
<input type="checkbox"/> <u>4796098</u>	January 1989	Giddings	386/47
<input type="checkbox"/> <u>4849978</u>	July 1989	Dishon et al.	395/182.04
<input type="checkbox"/> <u>5088081</u>	February 1992	Farr	369/54
<input type="checkbox"/> <u>5235601</u>	August 1993	Stallmo et al.	371/40.1
<input type="checkbox"/> <u>5303244</u>	April 1994	Watson	395/182.03
<input type="checkbox"/> <u>5404454</u>	April 1995	Parks	395/841

<input type="checkbox"/>	<u>5463758</u>	October 1995	Ottesen	395/441
<input type="checkbox"/>	<u>5471640</u>	November 1995	McBride	395/842
<input type="checkbox"/>	<u>5487160</u>	January 1996	Bemis	395/441
<input type="checkbox"/>	<u>5497478</u>	March 1996	Murata	395/484
<input type="checkbox"/>	<u>5499337</u>	March 1996	Gordon	395/182.04
<input type="checkbox"/>	<u>5519435</u>	May 1996	Anderson	395/441
<input type="checkbox"/>	<u>5519844</u>	May 1996	Stallmo	395/441

OTHER PUBLICATIONS

Third Intl Workshop, Nov. 1992, "The Design and Implementation of a Continuous Media Storage Server", P. Louher et al, pp. 69-80.

Proceedings of 26th Hawaii Intl Conf. on System Sciences, v. 1, 1993 IEEE "Disk Subsystem Load Balancing: Disk Striping vs. Conventional Data Placement", G. R. Ganger et al, pp. 40-49.

"Performance Analysis of a Dual Striping Strategy for Replicated Disk Arrays", by Merchant et al, 1993, pp. 148-157.

"A Synchronous Disk Interleaving", by Kim et al, 1991 IEEE, pp. 801-810.

"Analytic Modeling and Comparisons of striping Strategies for Replicated Disk Arrays", by Merchant et al, IEEE 1995, pp. 419-433.

"Chained Declustering", by Hsiao et al, IEEE 1990, pp. 456-465.

"Tuning of Striping Units in Disk--Array--Based File System", by Weikum et al, IEEE 1992, pp. 80-87.

"An Approach to Cost--Effective Terabyte Memory Systems", by Katz et al, IEEE 1992, pp. 395-400.

"A Performance Study of Three High Availability Data Replication Strategies" by Hsiao et al, IEEE 1991, pp. 18-28.

"Replicated Data Management in the Gamma Database Machine", by Hsiao et al, 1990 IEEE, pp. 79-84.

"Introduction to Redundant Arrays of Inexpensive Disks (RAID)", by Patterson et al, IEEE 1989, pp. 112-117.

"Disk Subsystem Load Balancing", by Ganger et al, IEEE 1993, pp. 40-49.

"Chained Declustering", by Golubchik et al, IEEE 1992, pp. 88-95.

"Systems Reliability and Availability Prediction", by Daya Perera, IEEE 1993, pp. 33-40.

"Communications--Intensive Workstations", by Katseff et al, IEEE 1992, pp. 24-28.

"LAN Lirchin", by Harbison, Robert, Feb. 1994, LAN Magazine, v9, n2, p. 93(14).

"Diamonds are Forever", by Radding, Alan, Midrance Systems, Feb. 9, 1993, v6, n3, p. 23(2).

ART-UNIT: 237

PRIMARY-EXAMINER: Meky; Moustafa M.

ATTY-AGENT-FIRM: Luludis; F. B.

ABSTRACT:

The reliability of supplying data stored in a plurality of different memories to different users is enhanced by (a) dividing each of the memories into primary and secondary sections, (b) partitioning the data into successive blocks and (c) storing the blocks of data in sequence in respective ones of the primary sections. Then storing in sequence the blocks of data that have been stored in the primary section of one of the memories in respective ones of the secondary sections of the other ones of said disks.

2 Claims, 8 Drawing figures

[Previous Doc](#)

[Next Doc](#)

[Go to Doc#](#)

h e b b g e e e f c e h

e ge

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#) [Generate Collection](#) [Print](#)

L34: Entry 149 of 153

File: USPT

Oct 14, 1997

DOCUMENT-IDENTIFIER: US 5678061 A

TITLE: Method for employing doubly striped mirroring of data and reassigning data streams scheduled to be supplied by failed disk to respective ones of remaining disksDrawing Description Text (5):

FIG. 3 illustrates a schedule that a storage node of FIG. 1 may create for the unloading of data blocks stored in an associated disk;

Detailed Description Text (3):

With that in mind, in an illustrative embodiment of the invention, multiplexer 75, FIG. 1, which may be, for example, an Asynchronous Transport Mode switch, receives data in the form of packets from storage nodes 50-1 through 50-N, respectively. Multiplexer 75 then routes the received packets over respective virtual circuit connections via communications paths 75-1 through 75-N to their intended destinations, in which a packet may carry a segment of a video program and in which the content of the video program may be doubly striped and mirrored in accord with the invention across storage disks 40-11 through 40-j4, as will be explained below in detail. Because of such striping, a storage node 50i may supply a packet carrying a respective portion of a segment of a particular video program to multiplexer 75 via a virtual channel assigned to the segments that are stored in storage disks 40k associated with storage node 50i and that are to be delivered to the same destination (subscriber), where N, j, i and k>1. Thus, segments of the content of a video program striped across the storage disks may be supplied to multiplexer 75 via respective ones of a number of different virtual data channels assigned for that purpose. Once such contents have been so delivered, then the assigned virtual data channels may be reassigned for some other but similar purpose (or left idle).

Detailed Description Text (5):

Processor 25, more particularly, first determines if it has the resources available to service the request. That is, if a disk is able to support n data streams and the content of a program is striped over N disks, then nN streams (users) can be supported from the array of N disks, where n and N>1. Thus, server 100 may service the user's request if the current number of data streams that are being supplied by the array of disks 40-11 through 40-j4 to multiplexer 75 is less than nd. Assuming that is the case, then processor 25 communicates with multiplexer 75 to obtain channel assignments that the storage nodes may use to sequentially transmit their respective video segments that form the requested video to multiplexer 75. Included in such communication is a request to establish a virtual connection between each of the assigned channels and a channel of one of the communications paths 76-1 through 76-N that will be used to route the program to the user. In addition, processor 25 establishes a schedule that the storage nodes 50i are to follow in the delivery of the segments to the user via multiplexer 75. Processor 25 then supplies the assigned channels and schedule to each of the storage nodes 50i via Local Area Network (LAN) 30, as discussed below. The schedule includes the identity of the storage node, e.g., 50-1, and associated disk, e.g., disk 40-11, storing the initial (first) segment of the requested program. The schedule also includes the time of day that the first segment of the program is to be unloaded from disk and supplied to multiplexer 75. Processor 25 then forms a message containing, inter

alia, the (a) schedule established for running the program, (b) identity of the program, (c) identity of the storage node containing the first segment of the program, (d) channel that has been assigned to the node whose address is contained in the header of the message, and (e) time of day that the first segment is to be transmitted to the user. Processor 25 then sends the message to the storage node identified in the message via LAN 30. Processor 25 then updates the message so that it is applicable to a next one of the storage nodes and sends the message to that node via LAN 30. In an illustrative embodiment of the invention, processor 25 sends the message first to the storage node 50i containing the first segment of the requested program. Processor 25 sequentially sends the message to the remaining storage nodes based on the order of their respective addresses and updates the message following each such transmission.

Detailed Description Text (6):

Since the storage nodes 50i are similar to one another, a discussion of one such node equally pertains to the other storage nodes. It is thus seen that storage node 50-1, FIG. 1, includes microprocessor 52-3 for communicating with host 25, in which the communications are typically directed to setting up a schedule for the delivery of respective video segments stored in buffer 52-2 to multiplexer 75 via the assigned data channel and for controlling the storing and unloading of the segments from buffer 52-2. Buffer 52-2 represents a dual buffer arrangement in which microprocessor 52-3 unloads segments of video from respective ones of the disks 40-11 through 40-14 and stores the unloaded segments in a first one of the dual buffers 52-2 during a first cycle. During that same cycle, e.g., a time period of one second, microprocessor 52-3 in turn unloads portions of respective segments of respective videos stored in the second one of the dual buffers 52-2 during a previous cycle. Adapter 52-1 reads a packet from the buffer and transmits the packet to multiplexer 75 via communication path 51-1 (which may be, e.g., optical fiber) and the channel assigned for that particular purpose. OC3 adapter 52-1 implements the well-known OC3 protocol for interfacing a data terminal, e.g., storage node 50-1, with an optical fiber communications path, e.g., path 51-1. SCSI adapter 52-4, on the other hand, implements a Small Computer System Interface between microprocessor 52-3 and its associated disks 40-11 through 40-14 via bus 45-1.

Detailed Description Text (9):

For the sake of simplicity and clarity, only disks 65-1 through 65-4 of a server 60, FIG. 2, are shown. (The other elements of the server 60, e.g., storage node, host processor, etc., are not shown.) Assume that the host processor has divided a video program into a plurality of sequential data blocks (segments) D0 through Di for storage in the disks, in accord with the invention. To do so, the host processor stores the data blocks D0 through Di in sequential order (i.e., round-robin order) in the primary sections P of disks 65-1 through 65-4, respectively, as shown in FIG. 2. That is, the host processor stripes the data blocks across the primary sections of disks 65-1 through 65-4. The host processor then makes a backup copy of the contents of the primary sections of the disks 65i, for example, 65-1. It does this, in accord with an aspect of the invention, by striping the data blocks forming such contents across the secondary sections of the other disks, i.e., disks 65-2 through 65-4, in round-robin order. For example, it is seen from FIG. 2 that the first three data blocks D0, D4 and D8 stored in disk 65-1 are also stored in the secondary (backup) sections S of disks 65-2 through 65-4, respectively. This is also the case for data blocks D12, D16 and D20 as well as the remaining contents of the primary section of disk 65-1. Similarly, the contents of the primary section of disk 65-2 are striped across disks 65-3, 65-4 and 65-1 in that order. For example, it is also seen from FIG. 2 that data blocks D1, D5 and D9 of the primary section of disk 65-2 are striped across disks 65-3, 65-4 and 65-1, respectively, and so on. Such backup striping is shown for the contents of the primary sections of disks 65-3 and 65-4.

Detailed Description Text (11):

h e b b g e e e f c e h

e g

Assume at this point that the host processor has received a request for the video program from a user and the host has communicated that request to each of the storage nodes 60-1 through 60-4 with the indication that the start of the program block D0 is to be delivered at time t0. Further assume that a cycle is one second and a block of data is three (3) megabits. Upon receipt of the request, node 60-1 generates a schedule for the unloading of the blocks of the program that are stored in associated disk 65-1 and stores the schedule in internal memory (not shown). The schedule so generated starts with the unloading of block D0 at time t0 minus 1 second for delivery to the associated server 300 multiplexer at time t0. The next entry, D4 is scheduled to be unloaded at time t0 plus 3 seconds for delivery at t0 plus 4 seconds. The next entry D8 is scheduled to be unloaded at time t0 plus 7 seconds for delivery at t0 plus 8 seconds, and so on. An illustrative example of such a schedule is shown in FIG. 3.

Detailed Description Text (12):

Storage node 60-2 generates a similar schedule with respect to data blocks D1, D5, D9, etc. That is, a storage node is programmed so that it determines its order in delivering the sequence of data blocks 35 forming the requested program to the associated multiplexer as a function of the identity of the node 60i having the first block of such data. Accordingly, the schedule that node 60-2 generates will indicate that data block D1, D5, D9, etc., are to be respectively delivered during cycles t0+1, t0+5, t0+9, and so on. Similarly, storage nodes 60-3 and 60-4 generate their own delivery schedules with respect to the data blocks of the requested program that are stored in their associated disks.

Detailed Description Text (16):

If the program exits block 405 via the 'no' path, then it increments (block 407) the number of data streams in batch i and sets a variable m to the address of node j. The program then assigns the user's request to batch i for transmission via bus m and the appropriate storage node 50m. The program (block 409) then sets C=(i+1) mod N and sets m=(m+1) mod N. The program (block 410) then determines if m=j and exits if that is the case. Otherwise, the program returns to block 408:

[Previous Doc](#)

[Next Doc](#)

[Go to Doc#](#)

[First Hit](#)[Fwd Refs](#)[Previous Doc](#)[Next Doc](#)[Go to Doc#](#)
 [Generate Collection](#) [Print](#)

L34: Entry 150 of 153

File: USPT

Jan 30, 1996

US-PAT-NO: 5488731

DOCUMENT-IDENTIFIER: US 5488731 A

TITLE: Synchronization method for loosely coupled arrays of redundant disk drives

DATE-ISSUED: January 30, 1996

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Mendelsohn; Noah	Arlington	MA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
International Business Machines Corporation	Armonk NY				02

APPL-NO: 08/ 352428 [PALM]

DATE FILED: December 8, 1994

PARENT-CASE:

This is a continuation of application Ser. No. 07/924,219 filed on Aug. 3, 1992, now abandoned.

INT-CL: [06] G06 F 11/00

US-CL-ISSUED: 395/800; 371/49.1, 371/47.1, 395/182.04, 395/182.12, 395/182.22

US-CL-CURRENT: 711/114; 714/14, 714/24, 714/6, 714/800

FIELD-OF-SEARCH: 395/800, 395/575, 364/DIG.1, 371/66, 371/51, 371/13, 371/49.1, 371/68.1

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<input type="checkbox"/> <u>T932005</u>	March 1975	Kruskal	235/153
<input type="checkbox"/> <u>4393500</u>	July 1983	Imazeki et al.	371/13
<input type="checkbox"/> <u>4410942</u>	October 1983	Milligan et al.	364/200
<input type="checkbox"/> <u>4419725</u>	December 1983	George et al.	364/200
<input type="checkbox"/> <u>4493083</u>	January 1985	Kinoshita	371/66

<input type="checkbox"/>	<u>4530054</u>	July 1985	Hamstra et al.	364/200
<input type="checkbox"/>	<u>4603406</u>	July 1986	Miyazaki et al.	365/229
<input type="checkbox"/>	<u>4654819</u>	March 1987	Stiffler et al.	364/900
<input type="checkbox"/>	<u>4697266</u>	September 1987	Finley	371/12
<input type="checkbox"/>	<u>4761785</u>	August 1988	Clark et al.	371/51
<input type="checkbox"/>	<u>4819159</u>	April 1989	Shipley et al.	364/DIG.1
<input type="checkbox"/>	<u>4942579</u>	July 1990	Goodlander et al.	371/51
<input type="checkbox"/>	<u>5043871</u>	August 1991	Nishigaki et al.	364/200
<input type="checkbox"/>	<u>5051887</u>	September 1991	Berger et al.	364/200
<input type="checkbox"/>	<u>5195100</u>	March 1993	Katz et al.	371/66
<input type="checkbox"/>	<u>5233618</u>	August 1993	Glider et al.	371/68.1
<input type="checkbox"/>	<u>5305326</u>	April 1994	Solomon et al.	371/11.1
<input type="checkbox"/>	<u>5313626</u>	May 1994	Jones et al.	395/575
<input type="checkbox"/>	<u>5379417</u>	January 1995	Lui et al.	395/575

FOREIGN PATENT DOCUMENTS

FOREIGN-PAT-NO	PUBN-DATE	COUNTRY	US-CL
WO90/06550	June 1990	WO	

OTHER PUBLICATIONS

Stonebraker et al., "Distributed Raid-A New Multiple Copy Algorithm", IEEE, 1990, pp. 430-437.
 Schloss et al., "Highly Redundant Management of Distributed Data", Workshop on the Management of Replicated Data, IEEE, Nov. 1990, pp. 91-95.
 D. Patterson et al., "A Case for Redundant Arrays of Inexpensive Disks (Raid)", ACM SIGMOD Conference, Chicago, Jun. 1988, pp. 109-116.

ART-UNIT: 232

PRIMARY-EXAMINER: Bowler; Alyssa H.

ASSISTANT-EXAMINER: Davis, Jr.; Walter D.

ATTY-AGENT-FIRM: Perman & Green

ABSTRACT:

A multiprocessor system includes a plurality of substantially identical nodes interconnected through a switching network, each node including a disk drive, NVRAM, and a processor. The system stores data in either a RAID or mirrored fashion across a plurality of disk drives in different nodes. When data is stored in a RAID arrangement, an NVRAM in a parity node is provided with an entry including the new data, a copy of old data from the node to which the new data is to be written, a copy of the old parity, and a synchronization state indicator. The parity node determines new parity and transmits the new data to the data node for storage. Upon receiving an acknowledgement, the parity node resets the synchronization indicator.

When power-up occurs, after a power failure, the parity node scans its NVRAM for any entry and upon finding one with a non-reset state indicator, transmits the new data to a destination data node for entry thereby synchronizing the contents of data and parity nodes. In a mirrored system, NVRAM in only one node has a data identifier entered into its NVRAM so that, upon a power failure and subsequent power-up, that entry enables the system to know which disk drives are in a non-synchronized state, and to cause actions that result in re-synchronization.

13 Claims, 6 Drawing figures

[Previous Doc](#)

[Next Doc](#)

[Go to Doc#](#)

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#) [Generate Collection](#) [Print](#)

L34: Entry 150 of 153

File: USPT

Jan 30, 1996

DOCUMENT-IDENTIFIER: US 5488731 A

TITLE: Synchronization method for loosely coupled arrays of redundant disk drives

Abstract Text (1):

A multiprocessor system includes a plurality of substantially identical nodes interconnected through a switching network, each node including a disk drive, NVRAM, and a processor. The system stores data in either a RAID or mirrored fashion across a plurality of disk drives in different nodes. When data is stored in a RAID arrangement, an NVRAM in a parity node is provided with an entry including the new data, a copy of old data from the node to which the new data is to be written, a copy of the old parity, and a synchronization state indicator. The parity node determines new parity and transmits the new data to the data node for storage. Upon receiving an acknowledgement, the parity node resets the synchronization indicator. When power-up occurs, after a power failure, the parity node scans its NVRAM for any entry and upon finding one with a non-reset state indicator, transmits the new data to a destination data node for entry thereby synchronizing the contents of data and parity nodes. In a mirrored system, NVRAM in only one node has a data identifier entered into its NVRAM so that, upon a power failure and subsequent power-up, that entry enables the system to know which disk drives are in a non-synchronized state, and to cause actions that result in re-synchronization.

Brief Summary Text (9):

The problem becomes more complex in architectures wherein multiple disks are written asynchronously, typically by separate controllers which can reside on separate processing nodes that are connected by a communication network. Where such disks are used for transaction processing systems, the prior art has made provision for using high level software transactions logs to enable resynchronization of the various disks, following a system failure.

Brief Summary Text (21):

A multiprocessor system includes a plurality of substantially identical nodes interconnected through a switching network, each node including a disk drive, NVRAM, and a processor. The system stores data in either a parity protected RAID or mirrored fashion across a plurality of disk drives in different nodes. When data is stored in a RAID arrangement, an NVRAM in a parity node is provided with an entry including the new data, a copy of the new parity, and a synchronization state indicator. The parity node determines new parity and transmits the new data to the data node for storage. Upon receiving an acknowledgement, the parity node resets the synchronization indicator. When power-up occurs, after a power failure, the parity node scans its NVRAM for any entry and upon finding one with a non-reset state indicator, transmits the new data to the destination data node for entry. In a mirrored system, NVRAM in only one node has a data identifier entered into its NVRAM so that, upon a power failure and subsequent power-up, that entry enables the system to know which disk drives are in a non-synchronized state.

Drawing Description Text (3):

FIG. 2 is a flow diagram indicating the procedures followed by the system of FIG. 1, in the case of mirrored-redundant data distribution.

Drawing Description Text (4):

h e b b g e e e f c e h

e g

FIG. 3 is a flow diagram of the procedure followed by the system of FIG. 1 subsequent to a power-up in a mirrored data redundancy arrangement.

Detailed Description Text (2):

Referring to FIG. 1, a multiprocessor system 10 comprises a plurality of nodes 12, each of which is substantially identical, all such nodes interconnected via a switch network 14. Each node 12 includes a disk drive 16, a processor 18, RAM 20 and an NVRAM 22. Processor 18, in the known manner, controls the operation disk drive 16, RAM 20, and NVRAM 22. The operation of system 10 is controlled by one or more nodal processors 18. The processor(s) may be located at a central controlling node, (e.g. node 24) or may be distributed throughout the nodal structure. Each node 12 must be accessible to a controlling node by means of switching network 14. Thus, any controlling node attempting to read or write a disk block must be in direct contact with all nodes in a parity group storing the block. In the alternative, the controlling node that attempts to read or write a disk block must be in contact with one of the disk nodes in the parity group, and the nodes in the parity group must be fully interconnected.

CLAIMS:

1. A multiprocessor system including a plurality of substantially identical nodes interconnected through a switch network, each node comprising disk drive means, nonvolatile random access memory (NVRAM) and a processor, said multiprocessor system storing RAID-structured data across disk drive means in a plurality of different nodes, said system performing a method comprising the steps of:
 - a. listing at least an identifier of a data segment to be updated by received update data in an NVRAM in a first node in response to a command to write said update data to said data segment;
 - b. sending said update data from said first node to a second node containing a copy of said data segment;
 - c. removing said listing of said identifier in said NVRAM in said first node only when said update data is written to disk drive means in said first node and after receiving a signal that said second node has recorded said update data;
 - d. causing each node, in the event of a power-up, to scan its NVRAM to find any listed identifiers of data segments contained therein; and
 - e. for any data segment denoted by a listed identifier in said NVRAM in said first node, causing a corresponding data segment in said second node to be in synchronism with said data segment denoted by said listed identifier in NVRAM in said first node.
7. A multiprocessor system including a plurality of substantially identical nodes interconnected through a switch network, each node comprising disk drive means, a nonvolatile random access memory (NVRAM) and a processor, said multiprocessor system storing RAID-structured data across disk drive means in a plurality of different nodes, said system performing a method comprising the steps of:

responding to a command to write new data to replace old data in a data segment stored in a first node, by storing in NVRAM in a different node which stores parity data corresponding to old data stored in said first node, an entry comprising said new data, a state indication, and a copy of new parity as calculated based upon an exclusive-or combination of old data from said first node, old parity from said different node and said new data;

transmitting said new data to said first node for storage therein, and upon receiving a signal acknowledging successful storage, causing said different-node to

reset said state indication; and

causing said different node, in the event of a power-up to scan its NVRAM for a said entry, and upon finding a said entry with a non-reset state indication, transmitting said new data to said first node.

12. A multiprocessor system including a plurality of substantially identical nodes interconnected through a switch network, each node comprising disk drive means, nonvolatile random access memory (NVRAM) and a processor, said multiprocessor system storing RAID-structured data across a disk drive means in a plurality of different nodes, said system comprising:

means for listing at least an identifier of a data segment to be updated by update data in an NVRAM in a first node in response to a command to write said update data to said data segment in said first node;

means for sending said update data from said first node to a second node containing a copy of said data segment;

means for removing said listing of said data segment in said NVRAM in said first node only when said update data is written to disk drive means in said first node and after receiving a signal that said second node has recorded said update data;

means for causing each node, in the event of a power-up, to scan its NVRAM to find any listed identifiers of data segments contained therein; and

means, responsive to finding a data segment identifier listed in said NVRAM in said first node, for causing a corresponding data segment in said second node to be in synchronism with said data segment listed in said NVRAM in said first node.

13. A multiprocessor system including a plurality of substantially identical nodes interconnected through a switch network, each node comprising disk drive means, a nonvolatile random access memory (NVRAM) and a processor, said multiprocessor system storing RAID-structured data across disk drive means in a plurality of different nodes, said system comprising:

means for responding to a command to write new data to replace old data in a data segment stored in a first node, by storing in NVRAM in a parity node which stores parity data corresponding to data stored in said first node, an entry comprising said new data, a state indication, and a copy of new parity as calculated based upon an exclusive-or combination of old data from said first node, old parity from said parity node and said new data;

means for transmitting said new data to said first node for storage therein, and upon receiving a signal acknowledging successful storage, causing said parity node to reset said state indication; and

means for causing said parity node, in the event of a power-up to scan its NVRAM for a said entry, and upon finding a said entry with a non-reset state indication, to transmit said new data to said first node.

[Previous Doc](#)

[Next Doc](#)

[Go to Doc#](#)

[First Hit](#) [Fwd Refs](#)[Previous Doc](#) [Next Doc](#) [Go to Doc#](#)
 [Generate Collection](#) [Print](#)

L19: Entry 60 of 68

File: USPT

Jun 23, 1998

US-PAT-NO: 5771354

DOCUMENT-IDENTIFIER: US 5771354 A

TITLE: Internet online backup system provides remote storage for customers using IDs and passwords which were interactively established when signing up for backup services

DATE-ISSUED: June 23, 1998

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Crawford; Christopher M.	Washington	DC	20016	

APPL-NO: 08/ 145825 [PALM]

DATE FILED: November 4, 1993

INT-CL: [06] G06 F 13/00

US-CL-ISSUED: 395/200.59; 395/200.49, 395/186, 395/620

US-CL-CURRENT: 709/229; 707/204, 709/219, 713/200

FIELD-OF-SEARCH: 364/200, 395/800, 395/876, 395/182.04, 395/182.09, 395/200.59, 395/200.49, 395/186, 395/620, 380/4

PRIOR-ART-DISCLOSED:

U. S. PATENT DOCUMENTS

 [Search Selected](#) [Search ALL](#) [Clear](#)

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<input type="checkbox"/> <u>4649479</u>	March 1987	Advani et al.	395/681
<input type="checkbox"/> <u>4901223</u>	February 1990	Rhyne	364/200
<input type="checkbox"/> <u>4954945</u>	September 1990	Inoue et al.	364/200
<input type="checkbox"/> <u>4982324</u>	January 1991	McConaughy et al.	395/200.52
<input type="checkbox"/> <u>4994963</u>	February 1991	Rorden et al.	364/200
<input type="checkbox"/> <u>5005122</u>	April 1991	Griffin et al.	395/200.01
<input type="checkbox"/> <u>5007082</u>	April 1991	Cummins	380/4
<input type="checkbox"/> <u>5023774</u>	June 1991	Sakai et al.	395/849
<input type="checkbox"/> <u>5089958</u>	February 1992	Horton	395/182.03
<input type="checkbox"/> <u>5107456</u>	April 1992	Schuur	395/310

<input type="checkbox"/>	<u>5109515</u>	April 1992	Laggis et al.	707/10
<input type="checkbox"/>	<u>5181238</u>	January 1993	Medamana et al.	379/95
<input type="checkbox"/>	<u>5210866</u>	May 1993	Milligan et al.	395/182.04
<input type="checkbox"/>	<u>5212784</u>	May 1993	Sparks	395/182.04
<input type="checkbox"/>	<u>5276867</u>	January 1994	Kenley et al.	707/104
<input type="checkbox"/>	<u>5317728</u>	May 1994	Tewis et al.	395/600
<input type="checkbox"/>	<u>5325430</u>	June 1994	Smyth et al.	380/4
<input type="checkbox"/>	<u>5353411</u>	October 1994	Nakaosa et al.	395/275
<input type="checkbox"/>	<u>5367698</u>	November 1994	Webber et al.	395/800
<input type="checkbox"/>	<u>5379418</u>	January 1995	Shimazaki et al.	395/182.09
<input type="checkbox"/>	<u>5404527</u>	April 1995	Irwin et al.	395/200.52
<input type="checkbox"/>	<u>5426594</u>	June 1995	Wright et al.	364/514
<input type="checkbox"/>	<u>5448718</u>	September 1995	Cohn et al.	395/404
<input type="checkbox"/>	<u>5497479</u>	March 1996	Hornbuckle	395/491
<input type="checkbox"/>	<u>5515502</u>	May 1996	Wood	395/182.13
<input type="checkbox"/>	<u>5535407</u>	July 1996	Yanagawa et al.	395/800
<input type="checkbox"/>	<u>5544320</u>	August 1996	Konrad	395/200.33
<input type="checkbox"/>	<u>5606719</u>	February 1997	Nichols et al.	395/876
<input type="checkbox"/>	<u>5696901</u>	December 1997	Konrad	395/200.33

OTHER PUBLICATIONS

Bershad, "Watchdogs: Extending the UNIX File System", Feb. 12, 1988, pp. 267-275.
 Needham, "Using Encryption for Authentication in Large Networks of Computers", Dec. 1978, pp. 993-999.
 Arneson, "Mass Storage Archiving In Network Enviroments", 1988, IEEE, pp. 45-50.
 To et al, "Knowledge Enginerring Approach To Data Centres Disaster Backup/Recovery Planning", 1989, IEEE, pp. 241-248.
 Ford et al, "Juggling features and Options in retrospect", Mar. 12, 1991, Macweek, p. 52.
 Steinberg, Jeffrey A., "Diverting Data from Disater (Advanced Systems Concepts Inc.'s Remote Shadow Option Disk Shadowing Software)," Software Review (Evaluation), Abstract.
 S. Armburst et al., PC Tech Journal, vol. 3, No. 9, "Forward Looking VDI," Sep. 1985, pp. 42-53.
 E. R. Carpenter, IBM Technical Disclosure Bulletin, vol. 34, No. 10B, "Single System Image With Network File System," IBM Corp., Mar. 1992, pp. 408-409.
 M. Tabes, Datamation, "Managing Storage Across the Enterprise," May 1992, pp. 36-40.
 P. Korzeniowski, Software Magazine, "Back to the Mainframe For Storage of LAN Data," Jul. 1992, pp. 73-77.
 Mark Ludwig, The Little Black Book of Computer Viruses, American Eagle publications, 1991, pp. 11-19, 23-29, 55-68, 69-76, 159-162.
 R. Waterman, Software Magazine, "Waiting for X500," Jul. 1992, 2 pages.
 R. Brown, Systems 3X/400, "Serving Many Masters," Mar. 1992, pp. 68-70.
 PC Computing, "Test and Evaluation Methods," Jun. 1992, pp. 238, 240, 247-248 and 262.
 Price list, "Tape Backup," 1 page.

Advertisement, "Let's Talk.RTM., Tired of the Runaround?," 1 page.
Advertisement, "PC/TCP.RTM. Plus," 1 page.
Advertisement, News 3X/400, "Mimix Hi-Net, Lakeview Technology Inc.," Sep. 1991, p. 140.
Advertisement, Network Innovation, "Server Based Backup," 1 page.
Advertisement, News 3X/400, "PC File Mail," Nov. 1991, p. 74.
Advertisement, "HandsOn Software, Monarch 36/400.TM. lets PC users get their hands on mid-range data," 1 page.
Advertisement, Andrew, "NetLynx.TM.: LAN Connectivity by the Numbers," 1 page.
Advertisement, "Innovation.RTM. Data Processing, FDR.Your DASD Management System," 1 page.
Advertisement, "Maynard.RTM., Maynard Delivers a Super Bundle," 1 page.
Advertisement, "Central Point Software Inc..RTM., Never Underestimate the Importance of Complete Protection," 1 page.
Advertisement, Business Communications, Gazelle, "Backup While You Work," Jun. 1992, p. 436.
Advertisement, "Works," 1 page.
Advertisement, Intel.RTM., Introducing fax software that can send, receive and run with your faxes, 1 page.
Advertisement, "LapLink, With LapLink you're only a phone call away," 1 page.
M. Tischer, PC Systems Programming, Abacus, Apr. 1993, pp. 157-171, 208-212, Appendix B and C.
Advertisement, PC/Computing, "SequeLink.TM., MiddleWare.TM.: Open Windows for Client/Server Computing," p. 427.
Advertisement, System Software Associates, Inc., BPCS Client/Server Application Software, 1 page.
Advertisement, "Rochester Software Connection, Inc., What's fast, fun and does the AS/400?" 1 page.
Robert Jourdan, Programers Problem Solver, Prentice Hall, 1986, pp. 369-394, 395-406.
The Waite Group, MS-DOS Developers Guide, Howard Sams & Co., 1989, pp. 281-354, 583, 739-748.
Advertisement, PC/Computing, Control up to 96 PC file servers with the 1 keyboard and monitor using . . . Commander by cybex, p. 433.
G. Campbell & J. Heim, PC World, "The Best of Online Services," Oct. 1993, pp. 225-240.
C. Shipley, Personal Computing, Ziff-Davis, "The World At Your Fingertips," Fall 1993, pp. 52-56.
J. Burgess, Washington Post Business Magazine Monday Section, Wash Tech, Cybertalk, "Avoiding Highway Robbery," Nov. 1993, pp. 19,26.
K. Neely, Midrange Computing, "Terminal Emulator or Network: A Management Guide," Sep. 1993, pp. 56-60.
P. Borsook, InfoWorld, "Enterprise Computing Internetworking," Feb. 1, 1993, pp. 48-49.
F. Derfler & S. Rigney, PC Magazine, vol. 10, No. 15, "From LAN to WAN, Bringing Your Networks Together," Sep. 10, 1991, pp. 106-118.
F. Derfler & S. Rigney, PC Magazine, vol. 10, No. 15, "Bridges and Routers, Smart Links Between LAN Segments," Sep. 10, 1991, pp. 121-169.
M. Thompson & S. Rigney PC Magazine, vol. 10, No. 15, "Access Servers, Dialing Up the LAN," Sep. 10, 1991, pp. 177-227.
J. Salemi, PC Magazine, vol. 10, No. 15, "Bulletin Board Software, Tools for Wide-Area Communications," Sep. 10, 1991, pp. 231-300.
M. Stone, PC Magazine, vol. 10, No. 15, "E-Mail LAN Links, Gateways to the World," Sep. 10, 1991, pp. 303-345.
F. Derfler & K. Maxwell, PC Magazine, vol. 10, No. 15, "Linking LANs, The Media Move The Message," Sep. 10, 1991, pp. 351-374.
F. Derfler & K. Maxwell, PC Magazine, vol. 10, No. 15, "Public Data Networks, Reliable Relays," Sep. 10, 1991, pp. 177-227.
P. Stephenson, LAN Magazine, "Chip Away at Viruses," Oct. 1992, pp. 95-102.
PC Computing, "Mobile Computing, No-Hassle Solutions for Keeping Files in Sync,"

Jul. 1993, pp. 218-219.

Advertisement, MicroCom (R), "Carbon Copy, #1 in Remote Control," 1 page.

Advertisement, Symantec, "The Norton pcAnywhere TM, Controlling your Desktop PC From A Portable Used To Require A Little Training," 1 page.

Microsoft Corporation, Microsoft MS-DOS(R) 6 Users Guide, Chapter 7 Features for Your Laptop Computer, Using Interlink to Connect Two Computers, Copyright 1993, pp. 161-166.

Symantec Corp., the Norton pcAnywhere Users Guide, Copyright 1990, 1991.

IBM Corporation, Application System/400.TM., System Introduction, GC21-9766-1, Copyright 1988, 1989.

IBM Corporation, Application System/400.TM., Programming: System Reference Summary, SC21-8104-1, Copyright 1988, 1989.

IBM Corporation, Application System/400.TM., Programming: Control Language Reference Volume 1, SC21-9775-1, Copyright 1988, 1989.

IBM Corporation, Application System/400.TM., Programming: Control Language Programmer's Guide, SC21-8077-1, Copyright 1988, 1989.

IBM Corporation, Application System/400.TM., Programming: Database Guide, SC21-9659-1, Copyright 1988, 1989.

IBM Corporation, Application System/400.TM., Programming: Data Management Guide, SC21-9658-1, Copyright 1988, 1989.

IBM Corporation, Application System/400.TM., Programming: Work Management Guide, SC21-8078-1, Copyright 1988, 1989.

IBM Corporation, Application System/400.TM., PC Support: DOS Planning and Installation Guide, SC21-8195-0, Copyright 1990.

IBM Corporation, Application System/400.TM., PC Support: DOS Operations Reference, SC21-8195-01, Copyright 1990.

IBM Corporation, Application System/400.TM., PC Support: DOS Users Guide, SC21-8199-0, Copyright 1990.

IBM Corporation, Application System/400.TM., PC Support: DOS and Os/2 Technical Reference, SC21-8091-2, Copyright 1988, 1990.

IBM Corporation, Application System/400.TM., PC Support: DOS and OS/2 Messages and Problem Analysis Guide, SC21-8093-2, Copyright 1988, 1990.

ART-UNIT: 276

PRIMARY-EXAMINER: Asta; Frank J.

ASSISTANT-EXAMINER: Luu; Le Hien

ATTY-AGENT-FIRM: Nixon & Vanderhye P.C.

ABSTRACT:

This invention makes it possible for a customer computer to connect to an online service provider computer by phone, Internet, or other method, pay a fee to said service provider, and obtain additional processing and storage resources for the customer's computer. The resources can take the form of virtual storage and processing capabilities. These capabilities give the customer computer what appears to be additional local processing power and/or additional local storage, this storage possibly including preloaded software and/or data.

The additional resources made available to the customer computer can be used either to enhance the customers' local needs (such as access to virtual storage for additional disk space, or access to a more powerful processor of similar type for program execution), or these additional resources can be used by the customer computer to support services on-line that otherwise would be unavailable, impractical, or unaffordable. Examples of services include software and information rental, sales, and release update services, anti-viral services, backup and recovery services, and diagnostic and repair services, to name a few.

27 Claims, 68 Drawing figures

[Previous Doc](#)

[Next Doc](#)

[Go to Doc#](#)

h e b b g e e e f c e

e ge

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.